

(19)



(11)

EP 2 016 476 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent:
10.07.2019 Bulletin 2019/28

(51) Int Cl.:
G06F 1/32 (2019.01)

(21) Application number: **07794750.5**

(86) International application number:
PCT/US2007/011326

(22) Date of filing: **10.05.2007**

(87) International publication number:
WO 2007/133646 (22.11.2007 Gazette 2007/47)

(54) ADAPTIVE STORAGE SYSTEM INCLUDING HARD DISK DRIVE WITH FLASH INTERFACE

ADAPTIVES SPEICHERSYSTEM MIT FESTPLATTENLAUFWERK MIT FLASH-SCHNITTSTELLE
SYSTÈME DE STOCKAGE ADAPTATIF COMPRENANT UN LECTEUR DE DISQUE DUR À
INTERFACE FLASH

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IS IT LI LT LU LV MC MT NL PL PT RO SE SI SK TR

• **YANG, Yun**
Los Altos, California 94024 (US)

(30) Priority: **10.05.2006 US 799151 P**
31.07.2006 US 820867 P
11.08.2006 US 503016

(74) Representative: **Grünecker Patent- und Rechtsanwälte**
PartG mbB
Leopoldstraße 4
80802 München (DE)

(43) Date of publication of application:
21.01.2009 Bulletin 2009/04

(56) References cited:
EP-A- 1 605 456 WO-A-2004/090889
US-A1- 2006 069 848 US-B1- 6 628 469

(73) Proprietor: **Marvell World Trade Ltd.**
St. Michael 14027 (BB)

• **"CompactFlash" [Online] 9 January 2006 (2006-01-09), WIKIPEDIA , XP002461645 Retrieved from the Internet: URL:http://web.archive.org/web/20060109003035/http://en.wikipedia.org/wiki/CompactFlash_II [retrieved on 2007-11-30] Paragraph "Microdrives"**

(72) Inventors:
• **SUTARDJA, Sehat**
Los Altos Hills, California 94022 (US)

EP 2 016 476 B1

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

FIELD OF THE INVENTION

[0001] The present invention relates to data storage systems, and more particularly to low power data storage systems.

BACKGROUND OF THE INVENTION

[0002] Laptop computers are powered using both line power and battery power. The processor, graphics processor, memory and display of the laptop computer consume a significant amount of power during operation. One significant limitation of laptop computers relates to the amount of time that the laptop can be operated using batteries without recharging. The relatively high power dissipation of the laptop computer usually corresponds to a relatively short battery life.

[0003] Referring now to FIG. 1A, an exemplary computer architecture 4 is shown to include a processor 6 with memory 7 such as cache. The processor 6 communicates with an input/output (I/O) interface 8. Volatile memory 9 such as random access memory (RAM) 10 and/or other suitable electronic data storage also communicates with the interface 8. A graphics processor 11 and memory 12 such as cache increase the speed of graphics processing and performance.

[0004] One or more I/O devices such as a keyboard 13 and a pointing device 14 (such as a mouse and/or other suitable device) communicate with the interface 8. A high power disk drive (HPDD) 15 such as a hard disk drive having one or more platters with a diameter greater than 1.8" provides nonvolatile memory, stores data and communicates with the interface 8. The HPDD 15 typically consumes a relatively high amount of power during operation. When operating on batteries, frequent use of the HPDD 15 will significantly decrease battery life. The computer architecture 4 also includes a display 16, an audio output device 17 such as audio speakers and/or other input/output devices that are generally identified at 18.

[0005] Referring now to FIG. 1B, an exemplary computer architecture 20 includes a processing chipset 22 and an I/O chipset 24. For example, the computer architecture may be a Northbridge/Southbridge architecture (with the processing chipset corresponding to the Northbridge chipset and the I/O chipset corresponding to the Southbridge chipset) or other similar architecture. The processing chipset 22 communicates with a processor 25 and a graphics processor 26 via a system bus 27. The processing chipset 22 controls interaction with volatile memory 28 (such as external DRAM or other memory), a Peripheral Component Interconnect (PCI) bus 30, and/or Level 2 cache 32. Level 1 cache 33 and 34 may be associated with the processor 25 and/or the graphics processor 26, respectively. Alternatively, an Accelerated Graphics Port (AGP) (not shown) communicates with the

processing chipset 22 instead of and/or in addition to the graphics processor 26. The processing chipset 22 is typically but not necessarily implemented using multiple chips. PCI slots 36 interface with the PCI bus 30.

[0006] The I/O chipset 24 manages the basic forms of input/output (I/O). The I/O chipset 24 communicates with an Universal Serial Bus (USB) 40, an audio device 41, a keyboard (KBD) and/or pointing device 42, and a Basic Input/Output System (BIOS) 43 via an Industry Standard Architecture (ISA) bus 44. Unlike the processing chipset 22, the I/O chipset 24 is typically (but not necessarily) implemented using a single chip, which is connected to the PCI bus 30. A HPDD 50 such as a hard disk drive also communicates with the I/O chipset 24. The HPDD 50 stores a full-featured operating system (OS) such as Windows XP® Windows 2000®, Linux and MAC®-based OS that is executed by the processor 25.

[0007] Prior art systems are described in documents EP-A-1 605 456 (MARVELL WORLD TRADE LTD [BB]), US 2006/069848 A1 (NALAWADI RAJEEV K [US] ET AL), "CompactFlash" [Online] 9 January 2006 (2006-01-09), WIKIPEDIA, URL:http://web.archive.org/web/20060109003035/http://en.wikipedia.org/wiki/CompactFlash_II , US-B1-6 628 469 (HOYT ROGER F [US]).

SUMMARY OF THE INVENTION

[0008] A disk drive system according to the invention is defined in claim 1. Preferred embodiments are defined by the dependent claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The present invention will become more fully understood from the detailed description and the accompanying drawings, wherein:

FIG. 1A and 1B illustrate exemplary computer architectures according to the prior art;

FIG. 2A illustrates a first exemplary computer architecture according to the present invention with a primary processor, a primary graphics processor, and primary volatile memory that operate during a high power mode and a secondary processor and a secondary graphics processor that communicate with the primary processor, that operate during a low power mode and that employ the primary volatile memory during the low power mode;

FIG. 2B illustrates a second exemplary computer architecture according to the present invention that is similar to FIG. 2A and that includes secondary volatile memory that is connected to the secondary processor and/or the secondary graphics processor; FIG. 2C illustrates a third exemplary computer architecture according to the present invention that is similar to FIG. 2A and that includes embedded volatile memory that is associated with the secondary proc-

essor and/or the secondary graphics processor;
 FIG. 3A illustrates a fourth exemplary architecture according to the present invention for a computer with a primary processor, a primary graphics processor, and primary volatile memory that operate during a high power mode and a secondary processor and a secondary graphics processor that communicate with a processing chipset, that operate during the low power mode and that employ the primary volatile memory during the low power mode;
 FIG. 3B illustrates a fifth exemplary computer architecture according to the present invention that is similar to FIG. 3A and that includes secondary volatile memory connected to the secondary processor and/or the secondary graphics processor;
 FIG. 3C illustrates a sixth exemplary computer architecture according to the present invention that is similar to FIG. 3A and that includes embedded volatile memory that is associated with the secondary processor and/or the secondary graphics processor;
 FIG. 4A illustrates a seventh exemplary architecture according to the present invention for a computer with a secondary processor and a secondary graphics processor that communicate with an I/O chipset, that operate during the low power mode and that employ the primary volatile memory during the low power mode;
 FIG. 4B illustrates an eighth exemplary computer architecture according to the present invention that is similar to FIG. 4A and that includes secondary volatile memory connected to the secondary processor and/or the secondary graphics processor;
 FIG. 4C illustrates a ninth exemplary computer architecture according to the present invention that is similar to FIG. 4A and that includes embedded volatile memory that is associated with the secondary processor and/or the secondary graphics processor;
 and
 FIG. 5 illustrates a caching hierarchy according to the present invention for the computer architectures of FIGs. 2A-4C;
 FIG. 6 is a functional block diagram of a drive control module that includes a least used block (LUB) module and that manages storage and transfer of data between the low-power disk drive (LPDD) and the high-power disk drive (HPDD);
 FIG. 7A is a flowchart illustrating steps that are performed by the drive control module of FIG. 6;
 FIG. 7B is a flowchart illustrating alternative steps that are performed by the drive control module of FIG. 6;
 FIGs. 7C and 7D are flowcharts illustrating alternative steps that are performed by the drive control module of FIG. 6;
 FIG. 8A illustrates a cache control module that includes an adaptive storage control module and that controls storage and transfer of data between the LPDD and HPDD;

FIG. 8B illustrates an operating system that includes an adaptive storage control module and that controls storage and transfer of data between the LPDD and the HPDD;

FIG. 8C illustrates a host control module that includes an adaptive storage control module and that controls storage and transfer of data between the LPDD and HPDD;

FIG. 9 illustrates steps performed by the adaptive storage control modules of FIGs. 8A-8C;

FIG. 10 is an exemplary table illustrating one method for determining the likelihood that a program or file will be used during the low power mode;

FIG. 11A illustrates a cache control module that includes a disk drive power reduction module;

FIG. 11B illustrates an operating system that includes a disk drive power reduction module;

FIG. 11C illustrates a host control module that includes a disk drive power reduction module;

FIG. 12 illustrates steps performed by the disk drive power reduction modules of FIGs. 11A-11C;

FIG. 13 illustrates a multi-disk drive system including a high-power disk drive (HPDD) and a lower power disk drive (LPDD);

FIGs. 14-17 illustrate other exemplary implementations of the multi-disk drive system of FIG. 13;

FIG. 18 illustrates the use of low power nonvolatile memory such as flash memory or a low power disk drive (LPDD) for increasing virtual memory of a computer;

FIGs. 19 and 20 illustrates steps performed by the operating system to allocate and use the virtual memory of FIG. 18;

FIG. 21 is a functional block diagram of a Redundant Array of Independent Disks (RAID) system according to the prior art;

FIG. 22A is a functional block diagram of an exemplary RAID system according to the present invention with a disk array including X HPDD and a disk array including Y LPDD;

FIG. 22B is a functional block diagram of the RAID system of FIG. 22A where X and Y are equal to Z;

FIG. 23A is a functional block diagram of another exemplary RAID system according to the present invention with a disk array including Y LPDD that communicates with a disk array including X HPDD;

FIG. 23B is a functional block diagram of the RAID system of FIG. 23A where X and Y are equal to Z;

FIG. 24A is a functional block diagram of still another exemplary RAID system according to the present invention with a disk array including X HPDD that communicate with a disk array including Y LPDD;

FIG. 24B is a functional block diagram of the RAID system of FIG. 24A where X and Y are equal to Z;

FIG. 25 is a functional block diagram of a network attachable storage (NAS) system according to the prior art;

FIG. 26 is a functional block diagram of a network

attachable storage (NAS) system according to the present invention that includes the RAID system of FIGs. 22A, 22B, 23A, 23B, 24A and/or 24B and/or a multi-drive system according to FIGs. 6-17; FIG. 27 is a functional block diagram of a disk drive controller incorporating a flash memory and disk drive interface controller; FIG. 28 is a functional block diagram of the interface controller of FIG. 27; FIG. 29 is a functional block diagram of a multi-disk drive system with a flash interface; and FIG. 30 is a flowchart illustrating steps performed by the multi-disk drive of FIG. 29.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0010] The following description of the preferred embodiment(s) is merely exemplary in nature and is in no way intended to limit the invention, its application, or uses. For purposes of clarity, the same reference numbers will be used in the drawings to identify similar elements. As used herein, the term module and/or device refers to an application specific integrated circuit (ASIC), an electronic circuit, a processor (shared, dedicated, or group) and memory that execute one or more software or firmware programs, a combinational logic circuit, and/or other suitable components that provide the described functionality.

[0011] As used herein, the term "high power mode" refers to active operation of the host processor and/or the primary graphics processor of the host device. The term "low power mode" refers to low-power hibernating modes, off modes, and/or non-responsive modes of the primary processor and/or primary graphics processor when a secondary processor and a secondary graphics processor are operable. An "off mode" refers to situations when both the primary and secondary processors are off.

[0012] The term "low power disk drive" or LPDD refers to disk drives and/or microdrives having one or more platters that have a diameter that is less than or equal to 1.8". The term "high power disk drive" or HPDD refers to hard disk drives having one or more platters that have a diameter that is greater than 1.8". LPDDs typically have lower storage capacities and dissipate less power than the HPDDs. The LPDDs are also rotated at a higher speed than the HPDDs. For example, rotational speeds of 10,000 - 20,000 RPM or greater can be achieved with LPDDs.

[0013] The term HDD with non-volatile memory interface (IF) refers to a hard disk drive that is connectable to a host device via a standard semiconductor memory interface of the host. For example, the semiconductor memory interface can be a flash interface.

[0014] The HDD with a non-volatile memory IF communicates with the host via the non-volatile memory interface using a non-volatile memory interface protocol. The non-volatile memory interface used by the host and

the HDD with non-volatile memory interface can include flash memory having a flash interface, NAND flash with a NAND flash interface or any other type of semiconductor memory interface. The HDD with a non-volatile memory IF can be a LPDD and/or a HPDD. The HDD with a non-volatile memory IF will be described further below in conjunction with FIGs. 27 and 28. Additional details relating to the operation of a HDD with a flash IF can be found in U.S. Patent US 7788427. In each of the implementations set forth below, the LPDD can be implemented using the HDD (implemented as a HPDD and/or LPDD) with a non-volatile memory IF. Alternately, the HDD with a non-volatile memory IF can be a LPDD and/or HPDD used in addition to the disclosed LPDD and/or HPDD.

[0015] The computer architecture according to the present invention includes the primary processor, the primary graphics processor, and the primary memory (as described in conjunction with FIGs. 1A and 1B), which operate during the high power mode. A secondary processor and a secondary graphics processor are operated during the low power mode. The secondary processor and the secondary graphics processor may be connected to various components of the computer, as will be described below. Primary volatile memory may be used by the secondary processor and the secondary graphics processor during the low power mode. Alternatively, secondary volatile memory, such as DRAM and/or embedded secondary volatile memory such as embedded DRAM can be used, as will be described below.

[0016] The primary processor and the primary graphics processor dissipate relatively high power when operating in the high power mode. The primary processor and the primary graphics processor execute a full-featured operating system (OS) that requires a relatively large amount of external memory. The primary processor and the primary graphics processor support high performance operation including complex computations and advanced graphics. The full-featured OS can be a Windows®-based OS such as Windows XP®, a Linux-based OS, a MAC®-based OS and the like. The full-featured OS is stored in the HPDD 15 and/or 50.

[0017] The secondary processor and the secondary graphics processor dissipate less power (than the primary processor and primary graphics processor) during the low power mode. The secondary processor and the secondary graphics processor operate a restricted-feature operating system (OS) that requires a relatively small amount of external volatile memory. The secondary processor and secondary graphics processor may also use the same OS as the primary processor. For example, a pared-down version of the full-featured OS may be used. The secondary processor and the secondary graphics processor support lower performance operation, a lower computation rate and less advanced graphics. For example, the restricted-feature OS can be Windows CE® or any other suitable restricted-feature OS. The restricted-feature OS is preferably stored in nonvolatile memory

such as flash memory, a HDD with a non-volatile memory IF, a HPDD and/or a LPDD. Preferably, the full-featured and restricted-feature OS share a common data format to reduce complexity.

[0018] The primary processor and/or the primary graphics processor preferably include transistors that are implemented using a fabrication process with a relatively small feature size. In one implementation, these transistors are implemented using an advanced CMOS fabrication process. Transistors implemented in the primary processor and/or primary graphics processor have relatively high standby leakage, relatively short channels and are sized for high speed. The primary processor and the primary graphics processor preferably employ predominantly dynamic logic. In other words, they cannot be shut down. The transistors are switched at a duty cycle that is less than approximately 20% and preferably less than approximately 10%, although other duty cycles may be used.

[0019] In contrast, the secondary processor and/or the secondary graphics processor preferably include transistors that are implemented with a fabrication process having larger feature sizes than the process used for the primary processor and/or primary graphics processor. In one implementation, these transistors are implemented using a regular CMOS fabrication process. The transistors implemented in the secondary processor and/or the secondary graphics processor have relatively low standby leakage, relatively long channels and are sized for low power dissipation. The secondary processor and the secondary graphics processor preferably employ predominantly static logic rather than dynamic logic. The transistors are switched at a duty cycle that is greater than 80% and preferably greater than 90%, although other duty cycles may be used.

[0020] The primary processor and the primary graphics processor dissipate relatively high power when operated in the high power mode. The secondary processor and the secondary graphics processor dissipate less power when operating in the low power mode. In the low power mode, however, the computer architecture is capable of supporting fewer features and computations and less complex graphics than when operating in the high power mode. As can be appreciated by skilled artisans, there are many ways of implementing the computer architecture according to the present invention. Therefore, skilled artisans will appreciate that the architectures that are described below in conjunction with FIGs. 2A-4C are merely exemplary in nature and are not limiting.

[0021] Referring now to FIG. 2A, a first exemplary computer architecture 60 is shown. The primary processor 6, the volatile memory 9 and the primary graphics processor 11 communicate with the interface 8 and support complex data and graphics processing during the high power mode. A secondary processor 62 and a secondary graphics processor 64 communicate with the interface 8 and support less complex data and graphics processing during the low power mode. Optional nonvolatile memory

65 such as a LPDD 66 and/or flash memory and/or a HDD with a non-volatile memory IF 69 communicates with the interface 8 and provides low power nonvolatile storage of data during the low power and/or high power modes. The HDD with a non-volatile memory IF can be a LPDD and/or a HPDD. The HPDD 15 provides high power/capacity nonvolatile memory. The nonvolatile memory 65 and/or the HPDD 15 are used to store the restricted feature OS and/or other data and files during the low power mode.

[0022] In this example, the secondary processor 62 and the secondary graphics processor 64 employ the volatile memory 9 (or primary memory) while operating in the low-power mode. To that end, at least part of the interface 8 is powered during the low power mode to support communications with the primary memory and/or communications between components that are powered during the low power mode. For example, the keyboard 13, the pointing device 14 and the primary display 16 may be powered and used during the low power mode. In all of the embodiments described in conjunction with FIGs. 2A-4C, a secondary display with reduced capabilities (such as a monochrome display) and/or a secondary input/output device can also be provided and used during the low power mode.

[0023] Referring now to FIG. 2B, a second exemplary computer architecture 70 that is similar to the architecture in FIG. 2A is shown. In this example, the secondary processor 62 and the secondary graphics processor 64 communicate with secondary volatile memory 74 and/or 76. The secondary volatile memory 74 and 76 can be DRAM or other suitable memory. During the low power mode, the secondary processor 62 and the secondary graphics processor 64 utilize the secondary volatile memory 74 and/or 76, respectively, in addition to and/or instead of the primary volatile memory 9 shown and described in FIG. 2A.

[0024] Referring now to FIG. 2C, a third exemplary computer architecture 80 that is similar to FIG. 2A is shown. The secondary processor 62 and/or secondary graphics processor 64 include embedded volatile memory 84 and 86, respectively. During the low power mode, the secondary processor 62 and the secondary graphics processor 64 utilize the embedded volatile memory 84 and/or 86, respectively, in addition to and/or instead of the primary volatile memory. In one embodiment, the embedded volatile memory 84 and 86 is embedded DRAM (eDRAM), although other types of embedded volatile memory can be used.

[0025] Referring now to FIG. 3A, a fourth exemplary computer architecture 100 according to the present invention is shown. The primary processor 25, the primary graphics processor 26, and the primary volatile memory 28 communicate with the processing chipset 22 and support complex data and graphics processing during the high power mode. A secondary processor 104 and a secondary graphics processor 108 support less complex data and graphics processing when the computer is in the

low power mode. In this example, the secondary processor 104 and the secondary graphics processor 108 employ the primary volatile memory 28 while operating in the low power mode. To that end, the processing chipset 22 may be fully and/or partially powered during the low power mode to facilitate communications therebetween. The HPDD 50 may be powered during the low power mode to provide high power volatile memory. Low power nonvolatile memory 109 (LPDD 110 and/or flash memory and/or HDD with a non-volatile memory IF 113) is connected to the processing chipset 22, the I/O chipset 24 or in another location and stores the restricted-feature operating system for the low power mode. The HDD with a non-volatile memory IF can be a LPDD and/or a HPDD.

[0026] The processing chipset 22 may be fully and/or partially powered to support operation of the HPDD 50, the LPDD 110, and/or other components that will be used during the low power mode. For example, the keyboard and/or pointing device 42 and the primary display may be used during the low power mode.

[0027] Referring now to FIG. 3B, a fifth exemplary computer architecture 150 that is similar to FIG. 3A is shown. Secondary volatile memory 154 and 158 is connected to the secondary processor 104 and/or secondary graphics processor 108, respectively. During the low power mode, the secondary processor 104 and the secondary graphics processor 108 utilize the secondary volatile memory 154 and 158, respectively, instead of and/or in addition to the primary volatile memory 28. The processing chipset 22 and the primary volatile memory 28 can be shut down during the low power mode if desired. The secondary volatile memory 154 and 158 can be DRAM or other suitable memory.

[0028] Referring now to FIG. 3C, a sixth exemplary computer architecture 170 that is similar to FIG. 3A is shown. The secondary processor 104 and/or secondary graphics processor 108 include embedded memory 174 and 176, respectively. During the low power mode, the secondary processor 104 and the secondary graphics processor 108 utilize the embedded memory 174 and 176, respectively, instead of and/or in addition to the primary volatile memory 28. In one embodiment, the embedded volatile memory 174 and 176 is embedded DRAM (eDRAM), although other types of embedded memory can be used.

[0029] Referring now to FIG. 4A, a seventh exemplary computer architecture 190 according to the present invention is shown. The secondary processor 104 and the secondary graphics processor 108 communicate with the I/O chipset 24 and employ the primary volatile memory 28 as volatile memory during the low power mode. The processing chipset 22 remains fully and/or partially powered to allow access to the primary volatile memory 28 during the low power mode.

[0030] Referring now to FIG. 4B, an eighth exemplary computer architecture 200 that is similar to FIG. 4A is shown. Secondary volatile memory 154 and 158 is connected to the secondary processor 104 and the second-

ary graphics processor 108, respectively, and is used instead of and/or in addition to the primary volatile memory 28 during the low power mode. The processing chipset 22 and the primary volatile memory 28 can be shut down during the low power mode.

[0031] Referring now to FIG. 4C, a ninth exemplary computer architecture 210 that is similar to FIG. 4A is shown. Embedded volatile memory 174 and 176 is provided for the secondary processor 104 and/or the secondary graphics processor 108, respectively in addition to and/or instead of the primary volatile memory 28. In this example, the processing chipset 22 and the primary volatile memory 28 can be shut down during the low power mode.

[0032] Referring now to FIG. 5, a caching hierarchy 250 for the computer architectures illustrated in FIGs. 2A-4C is shown. The HP nonvolatile memory HPDD 50 is located at a lowest level 254 of the caching hierarchy 250. Level 254 may or may not be used during the low power mode if the HPDD 50 is disabled and will be used if the HPDD 50 is enabled during the low power mode. The LP nonvolatile memory such as LPDD 110, flash memory and/or HDD with a non-volatile memory IF 113 is located at a next level 258 of the caching hierarchy 250. External volatile memory such as primary volatile memory, secondary volatile memory and/or secondary embedded memory is a next level 262 of the caching hierarchy 250, depending upon the configuration. Level 2 or secondary cache comprises a next level 266 of the caching hierarchy 250. Level 1 cache is a next level 268 of the caching hierarchy 250. The CPU (primary and/or secondary) is a last level 270 of the caching hierarchy. The primary and secondary graphics processor use a similar hierarchy.

[0033] The computer architecture according to the present invention provides a low power mode that supports less complex processing and graphics. As a result, the power dissipation of the computer can be reduced significantly. For laptop applications, battery life is extended.

[0034] Referring now to FIG. 6, a drive control module 300 or host control module for a multi-disk drive system includes a least used block (LUB) module 304, an adaptive storage module 306, and/or a LPDD maintenance module 308. The drive control module 300 controls storage and data transfer between a high-powered disk drive (HPDD) 310 such as a hard disk drive and a low-power disk drive (LPDD) 312 such as a microdrive based in part on LUB information. The drive control module 300 reduces power consumption by managing data storage and transfer between the HPDD and LPDD during the high and low power modes. As can be seen in FIG. 6, a HDD with a non-volatile memory IF 317 may be used as the LPDD and/or in addition to the LPDD. The drive control module 300 communicates with the HDD with a non-volatile memory IF 317 via a host non-volatile memory IF 315 and a host 313. The drive control module 300 may be integrated with the host 313 and/or the host non-vol-

atile memory IF 315.

[0035] The least used block module 304 keeps track of the least used block of data in the LPDD 312. During the low-power mode, the least used block module 304 identifies the least used block of data (such as files and/or programs) in the LPDD 312 so that it can be replaced when needed. Certain data blocks or files may be exempted from the least used block monitoring such as files that relate to the restricted-feature operating system only, blocks that are manually set to be stored in the LPDD 312, and/or other files and programs that are operated during the low power mode only. Still other criteria may be used to select data blocks to be overwritten, as will be described below.

[0036] During the low power mode during a data storing request the adaptive storage module 306 determines whether write data is more likely to be used before the least used blocks. The adaptive storage module 306 also determines whether read data is likely to be used only once during the low power mode during a data retrieval request. The LPDD maintenance module 308 transfers aged data from the LPDD to the HPDD during the high power mode and/or in other situations as will be described below.

[0037] Referring now to FIG. 7A, steps performed by the drive control module 300 are shown. Control begins in step 320. In step 324, the drive control module 300 determines whether there is a data storing request. If step 324 is true, the drive control module 300 determines whether there is sufficient space available on the LPDD 312 in step 328. If not, the drive control module 300 powers the HPDD 310 in step 330. In step 334, the drive control module 300 transfers the least used data block to the HPDD 310. In step 336, the drive control module 300 determines whether there is sufficient space available on the LPDD 312. If not, control loops to step 334. Otherwise, the drive control module 300 continues with step 340 and turns off the HPDD 310. In step 344, data to be stored (e.g. from the host) is transferred to the LPDD 312.

[0038] If step 324 is false, the drive control module 300 continues with step 350 and determines whether there is a data retrieving request. If not, control returns to step 324. Otherwise, control continues with step 354 and determines whether the data is located in the LPDD 312. If step 354 is true, the drive control module 300 retrieves the data from the LPDD 312 in step 356 and continues with step 324. Otherwise, the drive control module 300 powers the HPDD 310 in step 360. In step 364, the drive control module 300 determines whether there is sufficient space available on the LPDD 312 for the requested data. If not, the drive control module 300 transfers the least used data block to the HPDD 310 in step 366 and continues with step 364. When step 364 is true, the drive control module 300 transfers data to the LPDD 312 and retrieves data from the LPDD 312 in step 368. In step 370, control turns off the HPDD 310 when the transfer of the data to the LPDD 312 is complete.

[0039] Referring now to FIG. 7B, a modified approach that is similar to that shown in FIG. 7A is used and includes one or more adaptive steps performed by the adaptive storage module 306. When there is sufficient space available on the LPDD in step 328, control determines whether the data to be stored is likely to be used before the data in the least used block or blocks that are identified by the least used block module in step 372. If step 372 is false, the drive control module 300 stores the data on the HPDD in step 374 and control continues with step 324. By doing so, the power that is consumed to transfer the least used block(s) to the LPDD is saved. If step 372 is true, control continues with step 330 as described above with respect to FIG. 7A.

[0040] When step 354 is false during a data retrieval request, control continues with step 376 and determines whether data is likely to be used once. If step 376 is true, the drive control module 300 retrieves the data from the HPDD in step 378 and continues with step 324. By doing so, the power that would be consumed to transfer the data to the LPDD is saved. If step 376 is false, control continues with step 360. As can be appreciated, if the data is likely to be used once, there is no need to move the data to the LPDD. The power dissipation of the HPDD, however, cannot be avoided.

[0041] Referring now to FIG. 7C, a more simplified form of control can also be performed during low power operation. Maintenance steps can also be performed during high power and/or low power modes (using the LPDD maintenance module 308). In step 328, when there is sufficient space available on the LPDD, the data is transferred to the LPDD in step 344 and control returns to step 324. Otherwise, when step 328 is false, the data is stored on the HPDD in step 380 and control returns to step 324. As can be appreciated, the approach illustrated in FIG. 7C uses the LPDD when capacity is available and uses the HPDD when LPDD capacity is not available. Skilled artisans will appreciate that hybrid methods may be employed using various combinations of the steps of FIGs. 7A-7D.

[0042] In FIG. 7D, maintenance steps are performed by the drive control module 300 upon returning to the high power mode and/or at other times to delete unused or low use files that are stored on the LPDD. This maintenance step can also be performed in the low power mode, periodically during use, upon the occurrence of an event such as a disk full event, and/or in other situations. Control begins in step 390. In step 392, control determines whether the high power mode is in use. If not, control loops back to step 7D. If step 392 is true, control determines whether the last mode was the low power mode in step 394. If not, control returns to step 392. If step 394 is false, control performs maintenance such as moving aged or low use files from the LPDD to the HPDD in step 396. Adaptive decisions may also be made as to which files are likely to be used in the future, for example using criteria described above and below in conjunction with FIGs. 8A-10.

[0043] Referring now to FIGs. 8A and 8B, storage control systems 400-1, 400-2 and 400-3 are shown. In FIG. 8A, the storage control system 400-1 includes a cache control module 410 with an adaptive storage control module 414. The adaptive storage control module 414 monitors usage of files and/or programs to determine whether they are likely to be used in the low power mode or the high power mode. The cache control module 410 communicates with one or more data buses 416, which in turn, communicate with volatile memory 422 such as L1 cache, L2 cache, volatile RAM such as DRAM and/or other volatile electronic data storage. The buses 416 also communicate with low power nonvolatile memory 424 (such as flash memory, a HDD with a non-volatile memory IF and/or a LPDD) and/or high power nonvolatile memory 426 such as a HPDD 426. In FIG. 8B, a full-featured and/or restricted feature operating system 430 is shown to include the adaptive storage control module 414. Suitable interfaces and/or controllers (not shown) are located between the data bus and the HPDD and/or LPDD.

[0044] In FIG. 8C, a host control module 440 includes the adaptive storage control module 414. The host control module 440 communicates with a LPDD 426' and a hard disk drive 426'. The host control module 440 can be a drive control module, an Integrated Device Electronics (IDE), ATA, serial ATA (SATA) or other controller. As can be seen in FIG. 8C, a HDD with a non-volatile memory IF 431 may be used as the LPDD and/or in addition to the LPDD. The host control module 440 communicates with the HDD with a non-volatile memory IF 431 via a host non-volatile memory IF 429. The host control module 440 may be integrated with the host non-volatile memory IF 429.

[0045] Referring now to FIG. 9, steps performed by the storage control systems in FIGs. 8A-8C are shown. In FIG. 9, control begins with step 460. In step 462, control determines whether there is a request for data storage to nonvolatile memory. If not, control loops back to step 462. Otherwise, the adaptive storage control module 414 determines whether data is likely to be used in the low-power mode in step 464. If step 464 is false, data is stored in the HPDD in step 468. If step 464 is true, the data is stored in the nonvolatile memory 444 in step 474.

[0046] Referring now to FIG. 10, one way of determining whether a data block is likely to be used in the low-power mode is shown. A table 490 includes a data block descriptor field 492, a low-power counter field 493, a high-power counter field 494, a size field 495, a last use field 496 and/or a manual override field 497. When a particular program or file is used during the low-power or high-power modes, the counter field 493 and/or 494 is incremented. When data storage of the program or file is required to nonvolatile memory, the table 492 is accessed. A threshold percentage and/or count value may be used for evaluation. For example, if a file or program is used greater than 80 percent of the time in the low-power mode, the file may be stored in the low-power nonvolatile

memory such as flash memory, a HDD with a non-volatile memory IF and/or the microdrive. If the threshold is not met, the file or program is stored in the high-power non-volatile memory.

[0047] As can be appreciated, the counters can be reset periodically, after a predetermined number of samples (in other words to provide a rolling window), and/or using any other criteria. Furthermore, the likelihood may be weighted, otherwise modified, and/or replaced by the size field 495. In other words, as the file size grows, the required threshold may be increased because of the limited capacity of the LPDD.

[0048] Further modification of the likelihood of use decision may be made on the basis of the time since the file was last used as recorded by the last use field 496. A threshold date may be used and/or the time since last use may be used as one factor in the likelihood determination. While a table is shown in FIG. 10, one or more of the fields that are used may be stored in other locations and/or in other data structures. An algorithm and/or weighted sampling of two or more fields may be used.

[0049] Using the manual override field 497 allows a user and/or the operating system to manually override of the likelihood of use determination. For example, the manual override field may allow an L status for default storage in the LPDD, an H status for default storage in the HPDD and/or an A status for automatic storage decisions (as described above). Other manual override classifications may be defined. In addition to the above criteria, the current power level of the computer operating in the LPDD may be used to adjust the decision. Skilled artisans will appreciate that there are other methods for determining the likelihood that a file or program will be used in the high-power or low-power modes that fall within the teachings of the present invention.

[0050] Referring now to FIGs. 11A and 11B, drive power reduction systems 500-1, 500-2 and 500-3 (collectively 500) are shown. The drive power reduction system 500 bursts segments of a larger sequential access file such as but not limited audio and/or video files to the low power nonvolatile memory on a periodic or other basis. In FIG. 11A, the drive power reduction system 500-1 includes a cache control module 520 with a drive power reduction control module 522. The cache control module 520 communicates with one or more data buses 526, which in turn, communicate with volatile memory 530 such as L1 cache, L2 cache, volatile RAM such as DRAM and/or other volatile electronic data storage, nonvolatile memory 534 such as flash memory, a HDD with a non-volatile memory IF and/or a LPDD, and a HPDD 538. In FIG. 11B, the drive power reduction system 500-2 includes a full-featured and/or restricted feature operating system 542 with a drive power reduction control module 522. Suitable interfaces and/or controllers (not shown) are located between the data bus and the HPDD and/or LPDD.

[0051] In FIG. 11C, the drive power reduction system 500-3 includes a host control module 560 with an adaptive storage control module 522. The host control module

560 communicates with one or more data buses 564, which communicate with the LPDD 534' and the hard disk drive 538'. The host control module 560 can be a drive control module, an Integrated Device Electronics (IDE), ATA, serial ATA (SATA) and/or other controller or interface. As can be seen in FIG. 11C, a HDD with a non-volatile memory IF 531 may be used as the LPDD and/or in addition to the LPDD. The host control module 560 communicates with the HDD with a non-volatile memory IF 531 via a host non-volatile memory IF 529. The host control module 560 may be integrated with the host non-volatile memory IF 529.

[0052] Referring now to FIG. 12, steps performed by the drive power reduction systems 500 in FIGs. 11A-11C are shown. Control begins the step 582. In step 584, control determines whether the system is in a low-power mode. If not, control loops back to step 584. If step 586 is true, control continues with step 586 where control determines whether a large data block access is typically requested from the HPDD in step 586. If not, control loops back to step 584. If step 586 is true, control continues with step 590 and determines whether the data block is accessed sequentially. If not, control loops back to 584. If step 590 is true, control continues with step 594 and determines the playback length. In step 598, control determines a burst period and frequency for data transfer from the high power nonvolatile memory to the low power nonvolatile memory.

[0053] In one implementation, the burst period and frequency are optimized to reduce power consumption. The burst period and frequency are preferably based upon the spin-up time of the HPDD and/or the LPDD, the capacity of the nonvolatile memory, the playback rate, the spin-up and steady state power consumption of the HPDD and/or LPDD, and/or the playback length of the sequential data block.

[0054] For example, the high power nonvolatile memory is a HPDD that consumes 1-2W during operation, has a spin-up time of 4-10 seconds and a capacity that is typically greater than 20Gb. The low power nonvolatile memory is a microdrive that consumes 0.3-0.5W during operation, has a spin-up time of 1-3 seconds, and a capacity of 1-6Gb. As can be appreciated, the forgoing performance values and/or capacities will vary for other implementations. The HPDD may have a data transfer rate of 1Gb/s to the microdrive. The playback rate may be 10Mb/s (for example for video files). As can be appreciated, the burst period times the transfer rate of the HPDD should not exceed the capacity of the microdrive. The period between bursts should be greater than the spin-up time plus the burst period. Within these parameters, the power consumption of the system can be optimized. In the low power mode, if the HPDD is operated to play an entire video such as a movie, a significant amount of power is consumed. Using the method described above, the power dissipation can be reduced significantly by selectively transferring the data from the HPDD to the LPDD in multiple burst segments spaced at fixed intervals at a

very high rate (e.g., 100X the playback rate) and then the HPDD can be shut down. Power savings that are greater than 50% can easily be achieved.

[0055] Referring now to FIG. 13, a multi-disk drive system 640 according to the present invention is shown to include a drive control module 650 and one or more HPDD 644 and one or more LPDD 648. The drive control module 650 communicates with a host device via host control module 651. To the host, the multi-disk drive system 640 effectively operates the HPDD 644 and LPDD 648 as a unitary disk drive to reduce complexity, improve performance and decrease power consumption, as will be described below. The host control module 651 can be an IDE, ATA, SATA and/or other control module or interface.

[0056] Referring now to FIG. 14, in one implementation the drive control module 650 includes a hard disk controller (HDC) 653 that is used to control one or both of the LPDD and/or HPDD. A buffer 656 stores data that is associated with the control of the HPDD and/or LPDD and/or aggressively buffers data to/from the HPDD and/or LPDD to increase data transfer rates by optimizing data block sizes. A processor 657 performs processing that is related to the operation of the HPDD and/or LPDD.

[0057] The HPDD 648 includes one or more platters 652 having a magnetic coating that stores magnetic fields. The platters 652 are rotated by a spindle motor that is schematically shown at 654. Generally the spindle motor 654 rotates the platter 652 at a fixed speed during the read/write operations. One or more read/write arms 658 move relative to the platters 652 to read and/or write data to/from the platters 652. Since the HPDD 648 has larger platters than the LPDD, more power is required by the spindle motor 654 to spin-up the HPDD and to maintain the HPDD at speed. Usually, the spin-up time is higher for HPDD as well.

[0058] A read/write device 659 is located near a distal end of the read/write arm 658. The read/write device 659 includes a write element such as an inductor that generates a magnetic field. The read/write device 659 also includes a read element (such as a magneto-resistive (MR) element) that senses the magnetic field on the platter 652. A preamp circuit 660 amplifies analog read/write signals.

[0059] When reading data, the preamp circuit 660 amplifies low level signals from the read element and outputs the amplified signal to the read/write channel device. While writing data, a write current is generated that flows through the write element of the read/write device 659 and is switched to produce a magnetic field having a positive or negative polarity. The positive or negative polarity is stored by the platter 652 and is used to represent data. The LPDD 644 also includes one or more platters 662, a spindle motor 664, one or more read/write arms 668, a read/write device 669, and a preamp circuit 670.

[0060] The HDC 653 communicates with the host control module 651 and with a first spindle/voice coil motor (VCM) driver 672, a first read/write channel circuit 674,

a second spindle/VCM driver 676, and a second read/write channel circuit 678. The host control module 651 and the drive control module 650 can be implemented by a system on chip (SOC) 684. As can be appreciated, the spindle VCM drivers 672 and 676 and/or read/write channel circuits 674 and 678 can be combined. The spindleNCM drivers 672 and 676 control the spindle motors 654 and 664, which rotate the platters 652 and 662, respectively. The spindleNCM drivers 672 and 676 also generate control signals that position the read/write arms 658 and 668, respectively, for example using a voice coil actuator, a stepper motor or any other suitable actuator.

[0061] Referring now to FIGs. 15-17, other variations of the multi-disk drive system are shown. In FIG. 15, the drive control module 650 may include a direct interface 680 for providing an external connection to one or more LPDD 682. In one implementation, the direct interface is a Peripheral Component Interconnect (PCI) bus, a PCI Express (PCIX) bus, and/or any other suitable bus or interface.

[0062] In FIG. 16, the host control module 651 communicates with both the LPDD 644 and the HPDD 648. A low power drive control module 650LP and a high power disk drive control module 650HP communicate directly with, the host control module. Zero, one or both of the LP and/or the HP drive control modules can be implemented as a SOC. As can be seen in FIG. 16, a HDD with a non-volatile memory IF 695 may be used as the LPDD and/or in addition to the LPDD. The host control module 651 communicates with the HDD with a non-volatile memory IF 695 via a host non-volatile memory IF 693. The host control module 651 may be integrated with the host non-volatile memory IF 693.

[0063] In FIG. 17, one exemplary LPDD 682 is shown to include an interface 690 that supports communications with the direct interface 680. As set forth above, the interfaces 680 and 690 can be a Peripheral Component Interconnect (PCI) bus, a PCI Express (PCIX) bus, and/or any other suitable bus or interface. The LPDD 682 includes an HDC 692, a buffer 694 and/or a processor 696. The LPDD 682 also includes the spindleNCM driver 676, the read/write channel circuit 678, the platter 662, the spindle motor 665, the read/write arm 668, the read element 669, and the preamp 670, as described above. Alternately, the HDC 653, the buffer 656 and the processor 658 can be combined and used for both drives. Likewise the spindle/VCM driver and read channel circuits can optionally be combined. In the embodiments in FIGs. 13-17, aggressive buffering of the LPDD is used to increase performance. For example, the buffers are used to optimize data block sizes for optimum speed over host data buses.

[0064] In conventional computer systems, a paging file is a hidden file on the HPDD or HP nonvolatile memory that is used by the operating system to hold parts of programs and/or data files that do not fit in the volatile memory of the computer. The paging file and physical mem-

ory, or RAM, define virtual memory of the computer. The operating system transfers data from the paging file to memory as needed and returns data from the volatile memory to the paging file to make room for new data.

5 The paging file is also called a swap file.

[0065] Referring now to FIGs. 18-20, the present invention utilizes the LP nonvolatile memory such as the LPDD, a HDD with a non-volatile memory IF and/or flash memory to increase the virtual memory of the computer system. In FIG. 18, an operating system 700 allows a user to define virtual memory 702. During operation, the operating system 700 addresses the virtual memory 702 via one or more buses 704. The virtual memory 702 includes both volatile memory 708 and LP nonvolatile memory 710 such as flash memory, a HDD with a non-volatile memory IF and/or a LPDD.

[0066] Referring now to FIG. 19, the operating system allows a user to allocate some or all of the LP nonvolatile memory 710 as paging memory to increase virtual memory. In step 720, control begins. In step 724, the operating system determines whether additional paging memory is requested. If not, control loops back to step 724. Otherwise, the operating system allocates part of the LP non-volatile memory for paging file use to increase the virtual memory in step 728.

[0067] In FIG. 20, the operating system employs the additional LP nonvolatile memory as paging memory. Control begins in step 740. In step 744, control determines whether the operating system is requesting a data write operation. If true, control continues with step 748 and determines whether the capacity of the volatile memory is exceeded. If not, the volatile memory is used for the write operation in step 750. If step 748 is true, data is stored in the paging file in the LP nonvolatile memory in step 754. If step 744 is false, control continues with step 760 and determines whether a data read is requested. If false, control loops back to step 744. Otherwise, control determines whether the address corresponds to a RAM address in step 764. If step 764 is true, control reads data from the volatile memory in step 764 and continues with step 744. If step 764 is false, control reads data from the paging file in the LP nonvolatile memory in step 770 and control continues with step 744.

[0068] As can be appreciated, using LP nonvolatile memory such as flash memory, a HDD with a non-volatile memory IF and/or the LPDD to increase the size of virtual memory will increase the performance of the computer as compared to systems employing the HPDD. Furthermore, the power consumption will be lower than systems using the HPDD for the paging file. The HPDD requires additional spin-up time due to its increased size, which increases data access times as compared to the flash memory, which has no spin-up latency, and/or the LPDD or a LPDD HDD with a non-volatile memory IF, which has a shorter spin-up time and lower power dissipation.

[0069] Referring now to FIG. 21, a Redundant Array of Independent Disks (RAID) system 800 is shown to include one or more servers and/or clients 804 that com-

municate with a disk array 808. The one or more servers and/or clients 804 include a disk array controller 812 and/or an array management module 814. The disk array controller 812 and/or the array management module 814 receive data and perform logical to physical address mapping of the data to the disk array 808. The disk array typically includes a plurality of HPDD 816.

[0070] The multiple HPDDs 816 provide fault tolerance (redundancy) and/or improved data access rates. The RAID system 800 provides a method of accessing multiple individual HPDDs as if the disk array 808 is one large hard disk drive. Collectively, the disk array 808 may provide hundreds of Gb to 10's to 100's of Tb of data storage. Data is stored in various ways on the multiple HPDDs 816 to reduce the risk of losing all of the data if one drive fails and to improve data access time.

[0071] The method of storing the data on the HPDDs 816 is typically called a RAID level. There are various RAID levels including RAID level 0 or disk striping. In RAID level 0 systems, data is written in blocks across multiple drives to allow one drive to write or read a data block while the next is seeking the next block. The advantages of disk striping include the higher access rate and full utilization of the array capacity. The disadvantage is there is no fault tolerance. If one drive fails, the entire contents of the array become inaccessible.

[0072] RAID level 1 or disk mirroring provides redundancy by writing twice - once to each drive. If one drive fails, the other contains an exact duplicate of the data and the RAID system can switch to using the mirror drive with no lapse in user accessibility. The disadvantages include a lack of improvement in data access speed and higher cost due to the increased number of drives (2N) that are required. However, RAID level 1 provides the best protection of data since the array management software will simply direct all application requests to the surviving HPDDs when one of the HPDDs fails.

[0073] RAID level 3 stripes data across multiple drives with an additional drive dedicated to parity, for error Correction/recovery. RAID level 5 provides striping as well as parity for error recovery. In RAID level 5, the parity block is distributed among the drives of the array, which provides more balanced access load across the drives. The parity information is used to recovery data if one drive fails. The disadvantage is a relatively slow write cycle (2 reads and 2 writes are required for each block written). The array capacity is N-1, with a minimum of 3 drives required.

[0074] RAID level 0+1 involves striping and mirroring without parity. The advantages are fast data access (like RAID level 0), and single drive fault tolerance (like RAID level 1). RAID level 0+1 still requires twice the number of disks (like RAID level 1). As can be appreciated, there can be other RAID levels and/or methods for storing the data on the array 808.

[0075] Referring now to FIGs. 22A and 22B, a RAID system 834-1 according to the present invention includes a disk array 836 that includes X HPDD and a disk array

838 that includes Y LPDD. One or more clients and/or a servers 840 include a disk array controller 842 and/or an array management module 844. While separate devices 842 and 844 are shown, these devices can be integrated if desired. As can be appreciated, X is greater than or equal to 2 and Y is greater than or equal to 1. X can be greater than Y, less than Y and/or equal to Y. For example, FIG. 22B shows a RAID system 834-1' where X = Y = Z.

[0076] Referring now to FIGs. 23A, 23B, 24A and 24B, RAID systems 834-2 and 834-3 are shown. In FIG. 23A, the LPDD disk array 838 communicates with the servers/clients 840 and the HPDD disk array 836 communicates with the LPDD disk array 838. The RAID system 834-2 may include a management bypass path that selectively circumvents the LPDD disk array 838. As can be appreciated, X is greater than or equal to 2 and Y is greater than or equal to 1. X can be greater than Y, less than Y and/or equal to Y. For example, FIG. 23B shows a RAID system 834-2' where X = Y = Z. In FIG. 24A, the HPDD disk array 836 communicates with the servers/clients 840 and the LPDD disk array 838 communicates with the HPDD disk array 836. The RAID system 834-2 may include a management bypass path shown by dotted line 846 that selectively circumvents the LPDD disk array 838. As can be appreciated, X is greater than or equal to 2 and Y is greater than or equal to 1. X can be greater than Y, less than Y and/or equal to Y. For example, FIG. 24B shows a RAID system 834-3' where X = Y = Z. The strategy employed may include write through and/or write back in FIGs. 23A-24B.

[0077] The array management module 844 and/or the disk controller 842 utilizes the LPDD disk array 838 to reduce power consumption of the HPDD disk array 836. Typically, the HPDD disk array 808 in the conventional RAID system in FIG. 21 is kept on at all times during operation to support the required data access times. As can be appreciated, the HPDD disk array 808 dissipates a relatively high amount of power. Furthermore, since a large amount of data is stored in the HPDD disk array 808, the platters of the HPDDs are typically as large as possible, which requires higher capacity spindle motors and increases the data access times since the read/write arms move further on average.

[0078] According to the present invention, the techniques that are described above in conjunction with FIGs. 6-17 are selectively employed in the RAID system 834 as shown in FIG. 22B to reduce power consumption and data access times. While not shown in FIGs. 22A and 23A-24B, the other RAID systems according to the present invention may also use these techniques. In other words, the LUB module 304, adaptive storage module 306 and/or the LPDD maintenance module that are described in FIGs. 6 and 7A-7D are selectively implemented by the disk array controller 842 and/or the array management controller 844 to selectively store data on the LPDD disk array 838 to reduce power consumption and data access times. The adaptive storage control module 414

that is described in FIGs. 8A-8C, 9 and 10 may also be selectively implemented by the disk array controller 842 and/or the array management controller 844 to reduce power consumption and data access times. The drive power reduction module 522 that is described FIGs. 11A-11C and 12 may also be implemented by the disk array controller 842 and/or the array management controller 844 to reduce power consumption and data access times. Furthermore, the multi-drive systems and/or direct interfaces that are shown in FIGs. 13-17 may be implemented with one or more of the HPDD in the HPDD disk array 836 to increase functionality and to reduce power consumption and access times.

[0079] Referring now to FIG. 25, a network attached storage (NAS) system 850 according to the prior art is shown to include storage devices 854, storage requesters 858, a file server 862, and a communications system 866. The storage devices 854 typically include disc drives, RAID systems, tape drives, tape libraries, optical drives, jukeboxes, and any other storage devices to be shared. The storage devices 854 are preferably but not necessarily object oriented devices. The storage devices 854 may include an I/O interface for data storage and retrieval by the requesters 858. The requesters 858 typically include servers and/or clients that share and/or directly access the storage devices 854.

[0080] The file server 862 performs management and security functions such as request authentication and resource location. The storage devices 854 depend on the file server 862 for management direction, while the requesters 858 are relieved of storage management to the extent the file server 862 assumes that responsibility. In smaller systems, a dedicated file server may not be desirable. In this situation, a requester may take on the responsibility for overseeing the operation of the NAS system 850. As such, both the file server 862 and the requester 858 are shown to include management modules 870 and 872, respectively, though one or the other and/or both may be provided. The communications system 866 is the physical infrastructure through which components of the NAS system 850 communicate. It preferably has properties of both networks and channels, has the ability to connect all components in the networks and the low latency that is typically found in a channel.

[0081] When the NAS system 850 is powered up, the storage devices 854 identify themselves either to each other or to a common point of reference, such as the file server 862, one or more of the requesters 858 and/or to the communications system 866. The communications system 866 typically offers network management techniques to be used for this, which are accessible by connecting to a medium associated with the communications system. The storage devices 854 and requesters 858 log onto the medium. Any component wanting to determine the operating configuration can use medium services to identify all other components. From the file server 862, the requesters 858 learn of the existence of the storage devices 854 they could have access to, while the storage

devices 854 learn where to go when they need to locate another device or invoke a management service like backup. Similarly the file server 862 can learn of the existence of storage devices 854 from the medium services.

5 Depending on the security of a particular installation, a requester may be denied access to some equipment. From the set of accessible storage devices, it can then identify the files, databases, and free space available.

[0082] At the same time, each NAS component can identify to the file server 862 any special considerations it would like known. Any device level service attributes could be communicated once to the file server 862, where all other components could learn of them. For instance, a requester may wish to be informed of the introduction of additional storage subsequent to startup, this being triggered by an attribute set when the requester logs onto the file server 862. The file server 862 could do this automatically whenever new storage devices are added to the configuration, including conveying important characteristics, such as it being RAID 5, mirrored, and so on.

[0083] When a requester must open a file, it may be able to go directly to the storage devices 854 or it may have to go to the file server for permission and location information. To what extent the file server 854 controls access to storage is a function of the security requirements of the installation.

[0084] Referring now to FIG. 26, a network attached storage (NAS) system 900 according to the present invention is shown to include storage devices 904, requesters 908, a file server 912, and a communications system 916. The storage devices 904 include the RAID system 834 and/or multi-disk drive systems 930 described above in FIGs. 6-19. The storage devices 904 typically may also include disc drives, RAID systems, tape drives, tape libraries, optical drives, jukeboxes, and/or any other storage devices to be shared as described above. As can be appreciated, using the improved RAID systems and/or multi-disk drive systems 930 will reduce the power consumption and data access times of the NAS system 900.

[0085] Referring now to FIG. 27, a disk drive controller incorporating a non-volatile memory and disk drive interface controller. In other words, the HDD of FIG. 27 has a non-volatile memory interface (hereinafter called HDD with non-volatile memory interface (IF)). The device of FIG. 27 allows a HDD to be connected to an existing non-volatile memory interface (IF) of a host device to provide additional nonvolatile storage.

[0086] The disk drive controller 1100 communicates with a host 1102 and a disk drive 1104. The HDD with a non-volatile memory IF includes the disk drive controller 1100 and the disk drive 1104. The disk drive 1104 typically has an ATA, ATA-CE, or IDE type interface. Also coupled to the disk drive controller 1100 is an auxiliary non-volatile memory 1106, which stores firmware code for the disk drive controller. In this case, the host 1102, while shown as a single block, typically includes as relevant components an industry standard non-volatile memory slot (connector) of the type for connecting to

commercially available non-volatile memory devices, which in turn is connected to a standard non-volatile memory controller in the host. This slot typically conforms to one of the standard types, for instance, MMC (Multi Media Card), SD (Secure Data), SD/MMC which is a combination of SD and MMC, HS-MMC (High Speed-MMC), SD/HS-MMC which is a combination of SD and HS-MMC, and Memory Stick. This list is not limiting.

[0087] A typical application is a portable computer or consumer electronic device such as MP3 music player or cellular telephone handset that has one application processor that communicates with an embedded non-volatile memory through a non-volatile memory interface. The non-volatile memory interface may include a flash interface, a NAND flash interface and/or other suitable non-volatile semiconductor memory interfaces. In accordance with this disclosure, rather than a non-volatile semiconductor memory, a hard disk drive or other type of disk drive is provided replacing the non-volatile semiconductor memory and using its interface signals. The disclosed method provides a non-volatile memory-like interface for a disk drive, which makes it easier to incorporate a disk drive in such a host system which normally only accepts flash memory. One advantage of a disk drive over flash memory as a storage device is far greater storage capacity for a particular cost.

[0088] Only minimum changes in the host non-volatile memory controller firmware and software need be made to incorporate the disk drive using the disclosed interface controller. Also, minimum command overhead is provided. Advantageously, there is open-ended data transfer for any particular read or write operation, in terms of the number of logic blocks transferred between the host and the disk drive. Also, no sector count of the disk drive need be provided by the host.

[0089] In certain embodiments the disk drive 1104 may be a small form factor (SFF) hard disk drive, which typically has a physical size of 650 x 15 x 70 mm. A typical data transfer rate of such SSF hard disk drive is 25 megabytes per second.

[0090] The functions of the disk drive controller 1100 of FIG. 27 are further explained below. The disk drive controller 1100 includes an interface controller 1110, which presents to the host system 1102 as a flash memory controller with a 14-line bus. The interface controller 1110 also performs the functions of host command interpretation and data flow control between the host 1102 and a buffer manager 1112. The buffer manager circuit 1112 controls, via a memory controller 1116, the actual buffer (memory), which may be an SRAM or DRAM buffer 1118 that may be included as part of the same chip as interface controller 1100 or be on a separate chip. The buffer manager provides buffering features that are described further below.

[0091] The buffer manager 1112 is also connected to a processor Interface/Servo and ID-Less/Defect Manager (MPIF/SAIUDM) circuit 1122, which performs the functions of track format generation and defect management.

The MPIF/SAIUDM circuit 1122, in turn, connects to the Advanced High Performance Bus (AHB) 1126. Connected to the AHB bus 1126 is a line cache 1128, and a processor 1130; a Tightly Coupled Memory (TCM) 1134 is associated with the processor 1130. The processor 1130 may be implemented by an embedded processor or by a microprocessor. The purpose of the line cache 1128 is to reduce code execution latency. It may be coupled to an external flash memory 1106.

[0092] The remaining blocks in the disk drive controller 1100 perform functions to support a disk drive and include the servo controller 1140, the disk formatter and error correction circuit 1142, and the read channel circuitry 1144, which connects to the pre-amplification circuit in the disk drive 1104. The 14-line parallel bus with 8 lines (0-7) may carry the bi-directional in/out (I/O) data. The remaining lines may carry the commands CLE, ALE, /CE, /RE, /WE and R/B respectively.

[0093] Referring now to Figure 28, the interface controller of Figure 27 is shown in more detail. The interface controller 1110 includes a flash controller (flash_ctl) block 1150, a flash register (flash_reg) block 1152, a flash FIFO wrapper (flash_fifo_wrapper) block 1154, and a flash system synchronization (flash_sys_syn) block 1156.

[0094] The flash register block 1152 is used for register access. It stores commands programmed by the processor 1130 and the host 1102. A flash state machine (not shown) in the flash controller 1150 decodes the incoming command from the host 1102 and provides the controls for the disk drive controller 1100. The flash FIFO wrapper 1154 includes a FIFO, which may be implemented by a 32x32 bi-directional asynchronous FIFO. It generates data and control signals for transferring data to and receiving data from the buffer manager 1112 via the buffer manager interface (BM IF). The transfer direction of the FIFO may be controlled by the commands stored in the flash register 1152. The flash system synchronization block 1156 synchronizes control signals between the interface controller and the buffer manager interface. It also generates a counter clear pulse (clk2_clr) for the flash FIFO wrapper 1154.

[0095] The flash controller 1150 may control the interface signal lines to implement a random read of the LPDD. The flash controller 1150 may control the interface signal lines to implement a random write of the LPDD. The flash controller 1150 may control the interface signal lines to implement a sequential read of the LPDD and may control the interface signal lines to implement a sequential write of the LPDD. The flash controller 1150 may control the interface signal lines to implement a transfer of commands between the control module and the LPDD. The flash controller 1150 may map a set of LPDD commands to a corresponding set of flash memory commands.

[0096] The register memory 1152 communicates with the interface controller and a LPDD processor via a processor bus. The register memory 1152 stores commands programmed by the LPDD processor and the control

module. The flash controller 1150 may store read data from the LPDD in the buffer memory to compensate for differences in data transfer rates between the control module and the LPDD and may send a data ready signal to the control module to indicate there is data in the memory buffer.

[0097] The flash controller 1150 may store write data from the control module in the buffer memory to compensate for differences in data transfer rates between the control module and the LPDD. The flash controller 1150 may send a data ready signal to the control module to indicate there is data in the memory buffer.

[0098] Referring now to FIG. 29, a functional block diagram of an embodiment of a multi-disk drive system with a flash interface is shown generally at 1200. While the preceding discussion related to the use of one disk drive (such as the low power or high power disk drive) with a flash interface, multiple disk drives can be connected via the flash interface. More particularly, the multi-disk drive system with a flash interface 1200 include a host flash interface 1206 that communicates with a flash interface of a host 1202. The host flash interface 1202 operates as described above. A drive control module 1208 selectively operates zero, one or both of the HPDD 1220 and the LPDD 1222. Control techniques that are described above with respect to operation of low power and high power modes can be performed by the drive control module 1208. In some implementations, the host flash Interface 1206 senses a power mode of the host and/or receives information that identifies a power mode of the host 1202.

[0099] Referring now to FIG. 30, a flowchart illustrating steps performed by the multi-disk drive of FIG. 29 are shown. Control begins with step 1230. In step 1232, control determines whether the host is on. If step 1232 is true, control determines whether the host is in a high power mode in step 1234. If step 1234 is true, control powers up the LPDD 1222 and/or the HPDD 1220 as needed in step 1236. If step 1234 is false, control determines whether the host is in a low power mode in step 1238. If step 1238 is true, control powers down the HPDD and operates the LPDD as needed to conserve power in step 1240. Control continues from step 1238 (if false) and step 1240 with step 1232.

[0100] As can be appreciated, the HDDs with flash interfaces that are described above can use the multi-disk drive with flash interface as described above. Furthermore, any of the control techniques described above with respect to systems with LPDD and HPDD can be used in the multi-disk drive with flash interface shown in FIG. 29.

[0101] Those skilled in the art can now appreciate from the foregoing description that the broad teachings of the present invention can be implemented in a variety of forms. Therefore, while this invention has been described in connection with particular examples thereof, the scope of the invention as defined in the appended claims should not be so limited since other modifications will become

apparent to the skilled practitioner upon a study of the drawings, the specification and the following claims.

5 Claims

1. A disk drive system (1204) for a host device with a high power mode and a low power mode, the disk drive system comprising:

a low power disk drive (1222), LPDD;
a high power disk drive (1220), HPDD;
a nonvolatile semiconductor memory interface (1206); and

a control module (1208) configured to communicate with the nonvolatile semiconductor memory interface, with the LPDD and the HPDD, wherein during a storing request for write data in the low power mode, the control module is configured to determine whether there is sufficient space available on the LPDD for the write data and to transfer the write data to the LPDD if sufficient space on the LPDD is available and to store the write data on the HPDD if insufficient space on the LPDD is available, wherein the LPDD and the HPDD are configured to communicate with the host device via the nonvolatile semiconductor memory interface (1206), and

wherein the nonvolatile semiconductor memory interface comprises a processor (1130), a line cache (1128) connected with a flash memory (1106), a buffer (1118) configured to store the write data from the control module, a buffer manager (1112) configured to control the buffer, the nonvolatile semiconductor memory interface further comprising an interface controller (1110) configured to communicate with the buffer manager (1112), the nonvolatile semiconductor memory interface further comprising a circuit (1122) for track format generation and defect management, said circuit being connected with the buffer manager, wherein the interface controller comprises:

a flash register (1152) in communication with a processor bus (1126), the processor bus being connected i) between the line cache (1128) and the interface controller and ii) between the processor and the circuit (1122) for track format generation and defect management, wherein the flash register is configured to store commands programmed by the host device and the processor, the interface controller further comprising a flash first-in first-out wrapper, (1154) and
a flash controller (1150) configured to com-

- communicate with a host semiconductor memory interface of the host device, the flash register, and the flash first-in-first-out wrapper (1154), wherein the flash first-in-first-out wrapper is configured to generate control signals for transferring the buffer data to and receiving the buffer data from the buffer manager.
2. The disk drive system of claim 1, wherein the control module further includes a LPDD maintenance module that is adapted to transfer data files from the LPDD to the HPDD during the high power mode to increase available disk space on the LPDD, and wherein:
- the LPDD maintenance module is adapted to transfer the data files based on at least one of age, size and likelihood of future use in the low power mode,
- the HPDD includes one or more platters,
- the one or more platters have a diameter that is greater than 1.8" and wherein the LPDD includes one or more platters, and
- the one or more platters have a diameter that is less than or equal to 1.8".
3. The disk drive system of claim 1, wherein:
- the nonvolatile semiconductor memory interface is adapted to sense a power mode of the host device; and
- the disk drive system is adapted to power down the HPDD and operate the LPDD if the host device is in a low power mode.
4. The disk drive system of claim 1, wherein:
- the flash controller is adapted to decode an incoming command from the host device and implement at least one of a read and a write of the LPDD; the flash first-in-first-out wrapper is adapted to generate control signals for transferring the buffer data to and receiving the buffer data from the buffer manager via a buffer manager interface and based on the commands stored in the flash register;
- the interface controller comprises a flash system synchronization block; and
- the flash synchronization block is adapted to synchronize control signals between the interface controller and the buffer manager interface and generate a counter clear pulse for the flash first-in-first-out wrapper.
5. The disk drive system of claim 1, wherein: the non-volatile semiconductor memory interface is adapted

to communicate with a nonvolatile memory connector of the host device configured for at least one of a multi-media card and a memory stick.

Patentansprüche

1. Diskettenlaufwerkssystem (1204) für eine Host-Einrichtung mit einem Hochleistungsmodus und einem Niedrigleistungsmodus, wobei das Diskettenlaufwerkssystem aufweist:

ein Niedrigleistungsdiskettenlaufwerk (1222), LPDD;

ein Hochleistungsdiskettenlaufwerk (1220), HPDD;

eine nicht-flüchtige Halbleiterspeicherschnittstelle (1206); und

ein Steuermodul (1208), das ausgebildet ist, mit der nicht-flüchtigen Halbleiterspeicherschnittstelle, dem LPDD und dem HPDD zu kommunizieren, wobei während einer Speicheranforderung bezüglich Schreibdaten in den Niedrigleistungsmodus das Steuermodul ausgebildet ist zu ermitteln, ob ausreichend Platz auf dem LPDD für die Schreibdaten verfügbar ist, und ausgebildet ist, die Schreibdaten zu dem LPDD zu übertragen, wenn ausreichend Platz auf dem LPDD verfügbar ist, und die Schreibdaten auf dem HPDD zu speichern, wenn nicht ausreichend Platz auf dem LPDD verfügbar ist, wobei das LPDD und das HPDD ausgebildet sind, mit der Host-Einrichtung über die nicht-flüchtige Halbleiterspeicherschnittstelle (1206) zu kommunizieren, und wobei die nicht-flüchtige Halbleiterspeicherschnittstelle einen Prozessor (1130), einen Leitungs-Cache-Speicher (1128), der mit einem Flash-Speicher (1106) verbunden ist, einen Puffer, der zum Speichern der Schreibdaten aus dem Steuermodul ausgebildet ist, und einen Pufferverwalter (1112), der zum Steuern des Puffers ausgebildet ist, aufweist, wobei die nicht-flüchtige Halbleiterspeicherschnittstelle ferner eine Schnittstellensteuerung (1110) aufweist, die ausgebildet ist, mit dem Pufferverwalter (1112) zu kommunizieren, wobei die nicht-flüchtige Halbleiterspeicherschnittstelle ferner eine Schaltung (1122) zur Spurformaterzeugung und Fehlerverwaltung aufweist, wobei die Schaltung mit dem Pufferverwalter verbunden ist, wobei die Schnittstellensteuerung umfasst:

ein Flash-Register (1152) in Verbindung mit einem Prozessorbus (1126), wobei der Prozessorbus angeschlossen ist i) zwischen dem Leitungs-Cache-Speicher (1128) und der Schnittstellensteuerung und ii) zwi-

- schen dem Prozessor und der Schaltung (1122) zur Spurformaterzeugung und Fehlerverwaltung,
wobei das Flash-Register ausgebildet ist, Befehle, die von der Host-Einrichtung und dem Prozessor programmiert sind, zu speichern, wobei die Schnittstellensteuerung ferner eine Flash-Zuerst-Hinein-Zuerst-Heraus-Hülleinheit (1154) und eine Flash-Steuerung (1150) aufweist, die ausgebildet ist, mit einer Host-Halbleiterspeicherschnittstelle der Host-Einrichtung, dem Flash-Register und der Flash-Zuerst-Hinein-Zuerst-Heraus-Hülleinheit (1154) zu kommunizieren, wobei die Flash-Zuerst-Hinein-Zuerst-Heraus-Hülleinheit ausgebildet ist, Steuersignale zur Übertragung der Pufferdaten zu und zum Empfang der Pufferdaten aus dem Pufferverwalter zu erzeugen.
2. Diskettenlaufwerkssystem nach Anspruch 1, wobei das Steuermodul ferner ein LPDD-Wartungsmodul enthält, das ausgebildet ist, Datendateien während des Hochleistungsmodus von dem LPDD zu dem HPDD zu übertragen, um den verfügbaren Diskettenplatz auf dem LPDD zu vergrößern, und wobei: das LPDD-Wartungsmodul ausgebildet ist, die Datendateien auf der Grundlage eines Alters und/oder einer Größe und/oder einer wahrscheinlichen künftigen Nutzung in dem Niedrigleistungsmodus zu übertragen, wobei das HPDD eine oder mehrere Festplatten enthält, wobei die eine oder die mehreren Festplatten einen Durchmesser haben, der größer als 1,8 " ist, und wobei das LPDD eine oder mehrere Festplatten enthält, und die eine oder die mehreren Festplatten einen Durchmesser haben, der kleiner oder gleich 1,8 " ist.
3. Diskettenlaufwerkssystem nach Anspruch 1, wobei: die nicht-flüchtige Halbleiterspeicherschnittstelle ausgebildet ist, einen Leistungsmodus der Host-Einrichtung zu erkennen; und das Diskettenlaufwerkssystem ausgebildet ist, das HPDD abzuschalten und das LPDD zu betreiben, wenn die Host-Einrichtung in einem Niedrigleistungsmodus ist.
4. Diskettenlaufwerkssystem nach Anspruch 1, wobei: die Flash-Steuerung ausgebildet ist, einen eintreffenden Befehl aus der Host-Einrichtung zu decodieren und einen Lesevorgang und/oder einen Schreibvorgang des LPDD zu implementieren; wobei die Flash-Zuerst-Hinein-Zuerst-Heraus-Hülleinheit ausgebildet ist, Steuersignale zur Übertragung der Pufferdaten zu und zum Empfangen der Pufferdaten aus dem Pufferverwalter über eine Pufferverwalterschnittstelle und auf der Basis der in dem Flash-Register gespeicherten Befehle zu übertragen; wobei die Schnittstellensteuerung einen Flash-Systemsynchronisierblock umfasst; und der Flash-Systemsynchronisierblock ausgebildet ist, Steuersignale zwischen der Schnittstellensteuerung und der Pufferverwalterschnittstelle zu synchronisieren und einen Zählerrücksetzpuls für die Flash-Zuerst-Hinein-Zuerst-Heraus-Hülleinheit zu erzeugen.
5. Diskettenlaufwerkssystem nach Anspruch 1, wobei: die nicht-flüchtige Halbleiterspeicherschnittstelle ausgebildet ist, mit einem nicht-flüchtigen Speicherverbinder der Host-Einrichtung zu kommunizieren, die für eine Multimediakarte und/oder einen Speicherstift konfiguriert ist.

Revendications

1. Système de lecteur de disque (1204) pour un dispositif hôte avec un mode haute puissance et un mode basse puissance, le système de lecteur de disque comprenant:
- un lecteur de disque basse puissance (1222), LPDD;
un lecteur de disque haute puissance (1220), HPDD;
une interface de mémoire à semi-conducteur non volatile (1206); et
un module de commande (1208) configuré pour communiquer avec l'interface de mémoire à semi-conducteur non volatile, avec le LPDD et le HPDD,
dans lequel, pendant une demande de stockage pour des données d'écriture en mode basse puissance, le module de commande est configuré pour déterminer s'il y a suffisamment d'espace disponible sur le LPDD pour les données d'écriture et pour transférer les données d'écriture au LPDD si un espace suffisant sur le LPDD est disponible et pour stocker les données d'écriture sur le disque HPDD si l'espace disponible sur le disque LPDD est insuffisant,
dans lequel le LPDD et le HPDD sont configurés pour communiquer avec le dispositif hôte par le biais de l'interface de mémoire à semi-conducteur non volatile (1206), et
dans lequel l'interface de mémoire à semi-conducteur non volatile comprend un processeur (1130), une mémoire cache de ligne (1128) connectée à une mémoire flash (1106), une mémoire tampon (1118) configurée pour stocker les

données d'écriture provenant du module de commande, un gestionnaire de mémoire tampon (1112) configuré pour contrôler la mémoire tampon, l'interface de mémoire à semi-conducteur non volatile comprenant en outre un contrôleur d'interface (1110) configuré pour communiquer avec le gestionnaire de mémoire tampon (1112), l'interface de mémoire à semi-conducteur non volatile comprenant en outre un circuit (1122) pour suivre la génération de format et la gestion des défaillances, ledit circuit étant connecté au gestionnaire de mémoire tampon, dans lequel le contrôleur d'interface comprend:

un registre flash (1152) en communication avec un bus de processeur (1126), le bus de processeur étant connecté i) entre la mémoire cache de ligne (1128) et le contrôleur d'interface et ii) entre le processeur et circuit (1122) pour suivre la génération de format et la gestion des défaillances, dans lequel le registre flash est configuré pour stocker des commandes programmées par le dispositif hôte et le processeur, le contrôleur d'interface comprenant en outre une enveloppe flash de type premier entré premier sorti (1154), et

un contrôleur de flash (1150) configuré pour communiquer avec une interface mémoire hôte à semi-conducteur du dispositif hôte, le registre flash et l'enveloppe flash de type premier entré, premier sorti (1154), l'enveloppe flash de type premier entré premier sorti étant configurée pour générer des signaux de commande permettant de transférer les données de mémoire tampon vers le gestionnaire de mémoire-tampon et de recevoir les données de mémoire tampon en provenance du gestionnaire de mémoire-tampon.

2. Système d'unité de disque selon la revendication 1, dans lequel le module de commande comprend en outre un module de maintenance LPDD adapté pour transférer des fichiers de données du LPDD vers le HPDD en mode haute puissance afin d'augmenter l'espace disque disponible sur le LPDD, et dans lequel:

le module de maintenance LPDD est adapté pour transférer les fichiers de données en fonction d'au moins l'âge, la taille et la probabilité d'utilisation future en mode basse puissance; le HPDD comprend un ou plusieurs plateaux, le ou les plateaux ont un diamètre supérieur à 1,8 "et le LPDD comprend un ou plusieurs plateaux, et le ou les plateaux ont un diamètre inférieur ou

égal à 1,8".

3. Système de lecteur de disque selon la revendication 1, dans lequel:

l'interface de mémoire à semi-conducteur non volatile est adaptée pour détecter un mode de puissance du dispositif hôte; et le système de lecteur de disque est adapté pour éteindre le HPDD et mettre en marche le LPDD si le dispositif hôte est en mode basse puissance.

4. Système de lecteur de disque selon la revendication 1, dans lequel:

le contrôleur flash est adapté pour décoder une commande entrante en provenance du dispositif hôte et enclencher au moins une lecture ou une écriture du LPDD;

l'enveloppe flash de type premier entré premier sorti est adaptée pour générer des signaux de commande pour transférer les données de la mémoire tampon gestionnaire de mémoire tampon et pour recevoir les données de mémoire tampon en provenance du gestionnaire de mémoire tampon par le biais d'une interface de gestion de mémoire tampon et sur la base des commandes stockées dans le registre flash;

le contrôleur d'interface comprend un bloc de synchronisation de système flash; et

le bloc de synchronisation flash est adapté pour synchroniser les signaux de commande entre le contrôleur d'interface et l'interface du gestionnaire de mémoire tampon et générer une impulsion claire de compteur pour l'enveloppe flash de type premier entré premier sorti.

5. Système de lecteur de disque selon la revendication 1, dans lequel:

l'interface de mémoire à semi-conducteur non volatile est adaptée pour communiquer avec un connecteur de mémoire non volatile du dispositif hôte configuré pour au moins une carte multimédia et une clé USB.




FIG. 1A
Prior Art




FIG. 1B
Prior Art




FIG. 2A




FIG. 2B




FIG. 2C




FIG. 3A




FIG. 3B




FIG. 3C





FIG. 4A







FIG. 5



Caching Hierarchy




FIG. 6


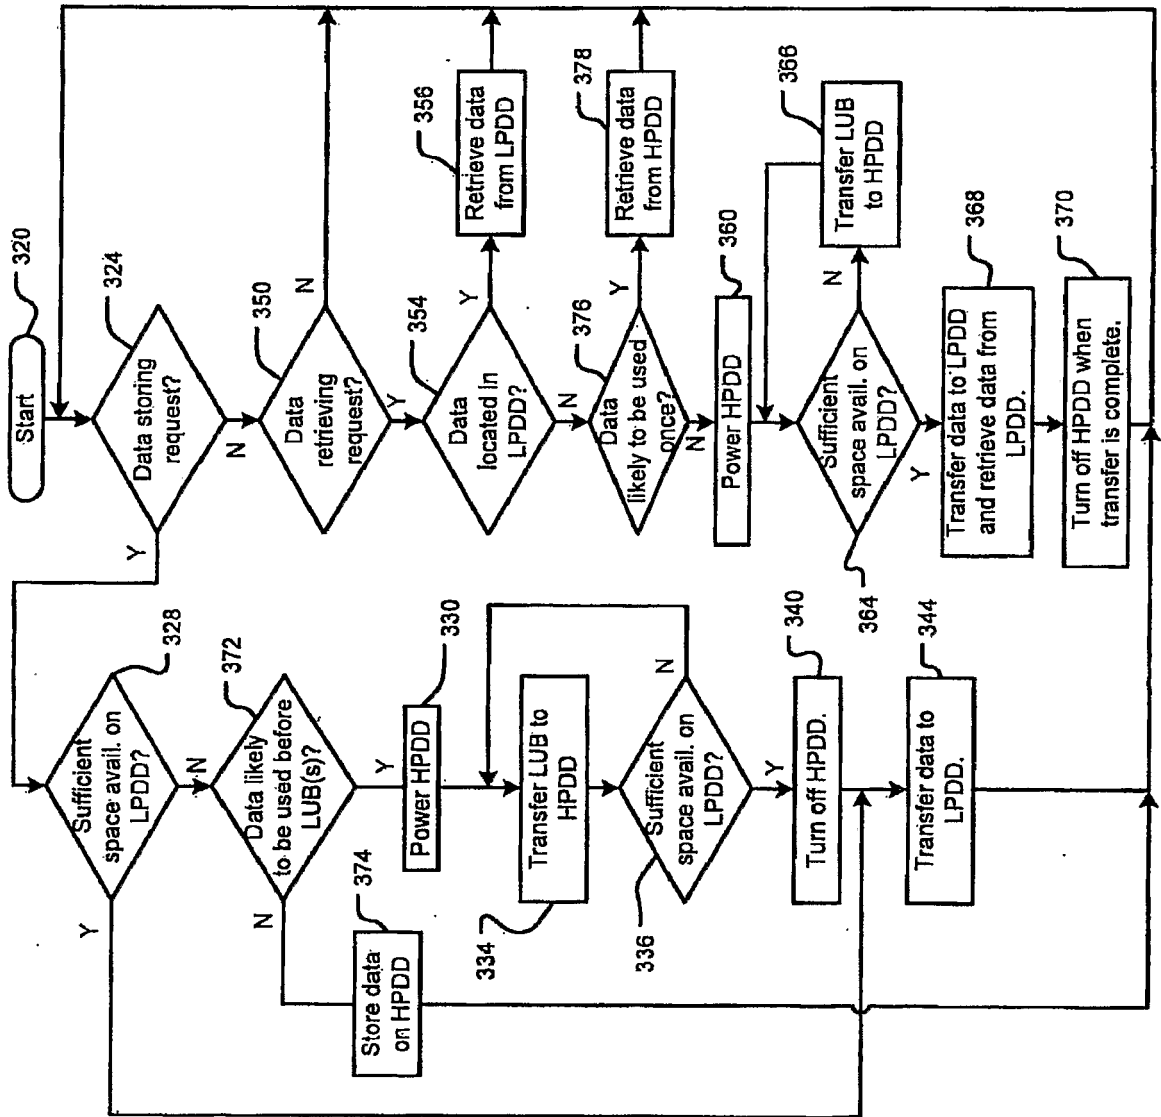


FIG. 7A

FIG. 7B






FIG. 7D

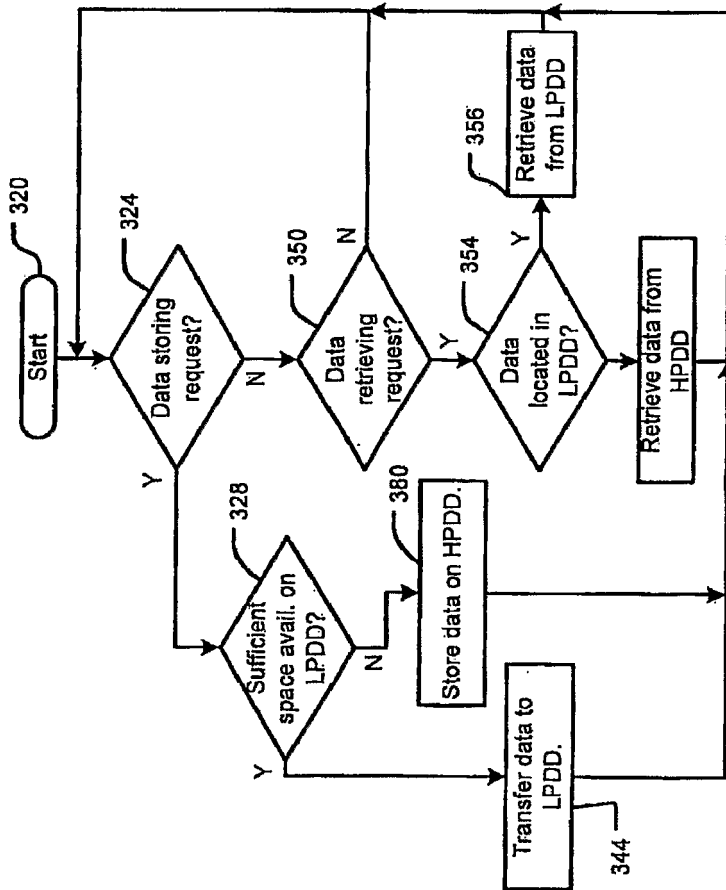


FIG. 7C




FIG. 8B




FIG. 8A




FIG. 8C




FIG. 9




FIG. 12

Data Block Descriptor	LP_Ctr	HP_Ctr	Size	LastUse	Manual
Program A	10	1	10M	D1	L
Program B	0	67	2G	D2	A
File A	0	45	80M	D3	H
File B	48	45	5G	D4	A

FIG. 10




FIG. 11B




FIG. 11A




FIG. 11C




FIG. 13




FIG. 14




FIG. 15




FIG. 16




FIG. 17




FIG. 18




FIG. 19

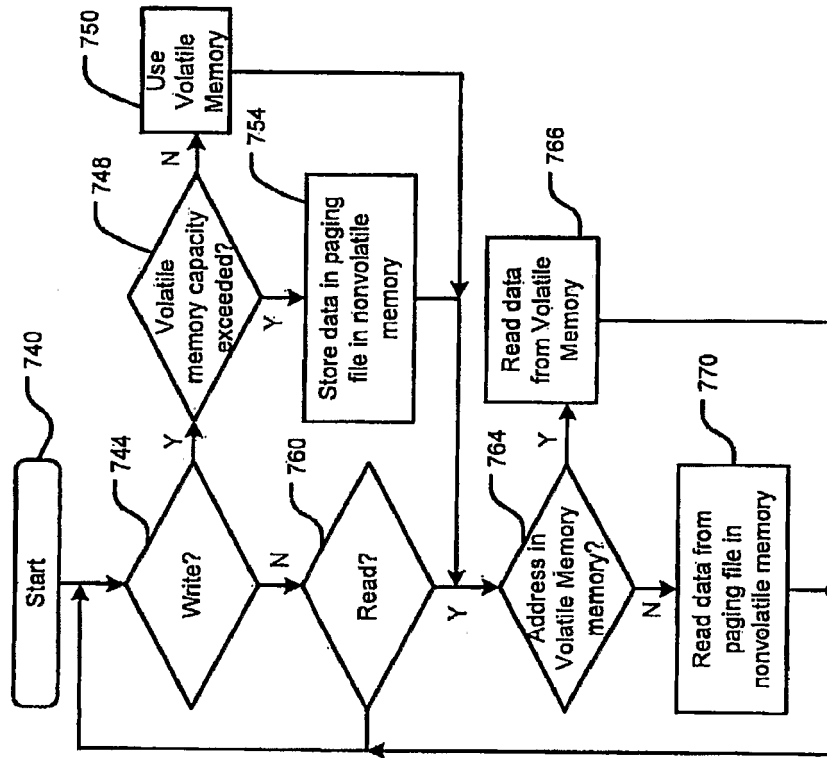


FIG. 20




FIG. 21
Prior Art




FIG. 22A




FIG. 22B




FIG. 23A




FIG. 24A




FIG. 23B




FIG. 24B




FIG. 25
Prior Art




FIG. 26




FIG. 27

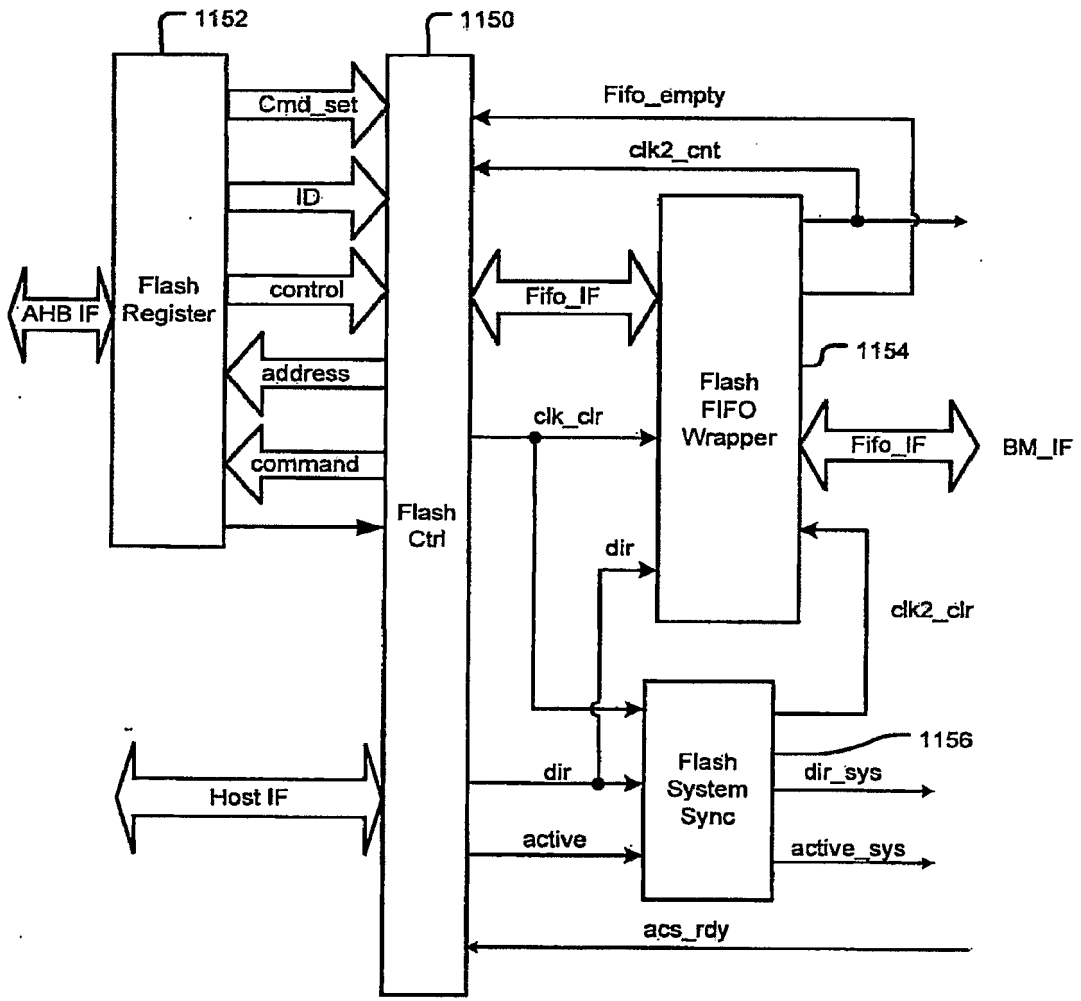


FIG. 28




FIG. 29




FIG. 30

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- EP 1605456 A [0007]
- US 2006069848 A1, NALAWADI RAJEEV K [0007]
- US 6628469 B1, HOYT ROGER F [0007]
- US 7788427 B [0014]

Non-patent literature cited in the description

- **WIKIPEDIA.** *CompactFlash*, 09 January 2006, http://web.archive.org/web/20060109003035/http://en.wikipedia.org/wiki/CompactFlash_II> [0007]